

SCALL: Software Component Allocator for Heterogeneous Embedded Systems

Ivan Švogar
University of Zagreb
Faculty of Organization and Informatics
Pavlinka 2, 42000 Varaždin, Croatia
isvogor@foi.hr

Jan Carlson
Mälardalen University
Mälardalen Real-Time Research Center
Box 883, 70123 Västerås, Sweden
jan.carlson@mdh.se

ABSTRACT

Allocation of software components on a heterogeneous computing platform involves hard decisions; handling different types of computing units with specific processing paradigms and a number of software components which require specific resources. An allocation method which explores the design space to provide a system architect with deployment alternatives in an early design phase can have a significant impact on the utilization of underlying platform. In this paper we present SCALL, an early prototype tool which uses heuristics and AHP for weighted multi-objective design space exploration to support systems architects in complex allocation decisions in early design phases.

Categories and Subject Descriptors

C.0 [Computer Systems Organization]: General—*System architectures*; D.2.2 [Software Engineering]: Design Tools and Techniques—*Decision tables*; D.2.8 [Software Engineering]: Metrics—*Performance measures*

General Terms

DESIGN, PERFORMANCE

Keywords

Software component allocation, heterogeneous, Eclipse, multi-objective, software architecture

1. INTRODUCTION

Heterogeneous computing has great potential in high-performance data processing, however with the benefits, it also brings new challenges for both hardware and software designers. These challenges are the subject of hardware–software codesign (HSCD), a field which investigates the concurrent design of hardware and software components in complex electronic systems. HSCD today seems to be skewed more towards hardware than software, however, with prevalence of off-the-shelf computing units this will change in future. A literature review [10] by Teich has shown that software engineering is overlapping with hardware design and that HSCD techniques should be known to both hardware and software engineers involved in complex electronic systems, e.g. networked systems, automotive, avionics, etc. Tool supported specification, modeling, partitioning and synthesis of subsystems is of utmost importance for this field in order to enable efficient design of complex systems with tight non-

functional constraints such as cost, performance, power, reliability, timing, etc. [12].

The task of allocation (i.e. placing, mapping) is complex since it involves several research fields besides HSCD, e.g. adaptive computing systems, performance optimization, component based software engineering, etc. and by no means is it straightforward. For example, a system with n components and m computing units leaves software architects with m^n possible allocations to choose from. This makes finding the best allocation laborious and even infeasible for large n and m . Therefore, a tool which supports and automates this process is essential for aiding in architectural decision making. In this paper we present SCALL — Software Component Allocator. It is an early prototype Eclipse plugin developed as part of our research on optimizing usage of heterogeneous computing platform by allocating software components.

2. OBJECTIVE

In order to address issues described in the introduction we are using a component based approach which efficiently deals with static and dynamic configurations, variant explosions, scalability, reusability, etc. This results in a system composed of independent, more manageable, reusable and replaceable building blocks which enable various allocations based on the specification of extra-functional properties (energy consumption, processing time, memory requirements, etc.). Since different allocations of software components result in different system performance [8], our research goal is to construct a framework which optimizes the allocation of software components on a heterogeneous computing platform with respect to specified extra-functional requirements. Although similar tools exist, they largely fail to address heterogeneous platforms [1, 5, 3].

3. FUNDAMENTALS: THE MATHEMATICAL MODEL

The basic elements of the (mathematical) model used by SCALL, are computing units, software components and the communication channels. The input for this model are matrices which quantify all the parameters necessary for making the allocation decision [11]. The assumptions for the model are the following: a) the system is built out of atomic components, i.e. a component cannot be distributed over several computing units, b) each component can be deployed to any computing unit or a platform specific equivalent exist.

The system is a subject to a set of constraints which spec-

ify whether a particular component is supposed to be (or not be) allocated to a particular computing unit. To represent that kind of system the mathematical model uses five matrices: \mathcal{T} , \mathcal{R} , \mathcal{C} , \mathcal{K} and \mathcal{B} ¹.

\mathcal{T} is a *resource consumption matrix* which specifies the resources required (i.e. consumed) by a component, since each component requires a certain amount of resources, which should be provided by the computing unit. This amount depends both on the kind of computing unit and the component. On the other hand, to specify resources which each computing unit can provide (e.g. available processing time, static memory, dynamic memory, energy etc.) we use the *computing unit resource matrix* \mathcal{R} .

The model considers communication in two aspects; a) hardware and b) software communication (as suggested in [6]). In one hand, hardware communication is limited by a physical constraint; bandwidth, represented by the *bandwidth matrix* \mathcal{B} . Additionally, since different computing units can use different communication channels and data types our model provides a *platform communication cost matrix* \mathcal{C} which is a placeholder for expressing additional processing due to heterogeneous nature of such systems.

On the other hand software communication can also cause overall system performance. Components which communicate intensively across computing units with high communication cost will have a larger impact on overall performance than ones communicating sporadically with lower data exchange. To express this constraint we define the *communication intensity matrix* \mathcal{K} .

Finally, function $p_i : S \rightarrow H$ is a *component mapping function* where $p_i = (p_1, \dots, p_n) \in \mathcal{P}$ defines a particular allocation (from the set of possible allocations \mathcal{P}) of components (S) to computing units (H), with $i = 1, \dots, m^n$. Elements of vector p_i represent one mapping, i.e. allocation of components to the computing units. To find the best one, allocations are evaluated and compared using the cost function (1):

$$w = \left(\sum_{k=1}^l f_k \sum_{i=1}^n t_{ip_k} + f_c \sum_{i \leq j} k_{ij} \cdot c_{p_i p_j} \right) \cdot \rho \cdot \kappa, \quad (1)$$

where: l is the number of resources used in matrix \mathcal{T} , n is the number of components, t is an element of a resource consumption matrix \mathcal{T} , k is an element of a communication intensity matrix \mathcal{K} , b is an element of the bandwidth matrix \mathcal{B} , c is an element of the communication cost matrix \mathcal{C} , f is an element of a trade-off vector F . Given an architect's preference about importance of different resources from matrix \mathcal{R} , a trade-off vector F is calculated and it provides weights for different resources. To calculate F we are using Analytic Hierarchy Process (AHP [7]). The multiplier ρ and κ are used to remove infeasible allocations. $\rho = 0$ when the allocation uses more resources than it is available, effectively disregarding it. Similarly, multiplier κ prevents an allocation in which the communication exceeds the available bandwidth. In all other cases both parameters are set to 1. To find the optimal, i.e. best, component allocation one needs to compare all the feasible solutions p_i from \mathcal{P} and select the one with the smallest $w > 0$. Having many components (n) and computing units (m), the search space

¹Complete and precise mathematical definitions for all the matrices can be found in [11].

$|\mathcal{P}| = m^n$ tends to be a large number often hard (infeasible) to find, so SCALL incorporates a framework to reduce the search space and find sub-optimal good enough solution.

4. SCALL: SOFTWARE COMPONENT ALLOCATOR TOOL

To implement SCALL, Eclipse was used, along with Eclipse Modeling Framework (EMF) and Graphical Modeling Project (GMF). This tool chain delivers a good platform for developing customized tools; EMF provides a modeling framework and code generation facilities, while GMF provides a set of generative components which create an infrastructure for development of graphical editors (Figure 2).

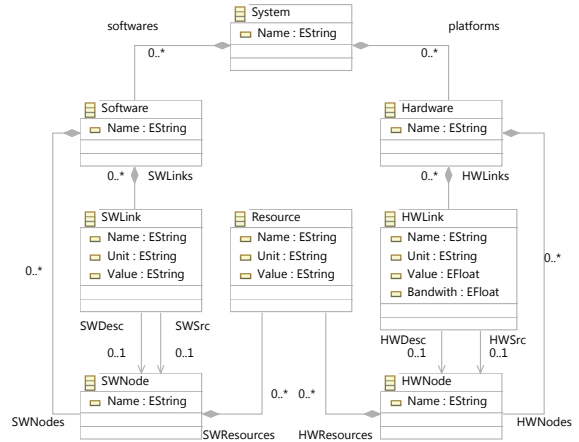


Figure 1: SCALL metamodel in Ecore notation

SCALL consists of two main parts; a) *Eclipse based model editor* and b) *PyAllocator*.

a) *Eclipse based model editor* – every model represented in SCALL editor corresponds to the metamodel depicted in Figure 1. The metamodel uses a standard EMF Ecore notation. Unlike UML or its profiles (e.g. Marte) which could have been also used for SCALL, our metamodel uses a small number of straightforward concepts. Having a custom model, we can focus on the recognition of important concepts without a burden of large industry standard models and reason about them more effectively. In future, if they prove to be valuable, these concepts can be added to standard models.

Eclipse based model editor is used to design the system. Figure 1 shows that the **System** consists of two main compartments: **Software** and **Hardware**. The **Software** is used to host a component model. It can consist of three main concepts; **SWNodes**, **SWLinks** and **Resources**. **SWNodes** represent software components which communicate via **SWLinks** and require some **Resources** to meet their execution specification. A **SWNode** can be associated with any number of links and resources, however all the **SWNodes** need to define values for all resources used in the model.

Similarly, the computing platform is represented by **Hardware** which can host **HWNodes**, **HWLinks** and **Resources**. **HWNodes** represent (heterogeneous) computing units which communicate via **HWLinks** and provide a set of **Resources**. **HWLinks** represent physical communication media for which the current metamodel recognizes two attributes: bandwidth and

communication cost (matrices B and C). **Software** compartment hosts **SWNodes** with corresponding resource demand. For **SWLinks** through which components communicate, allocation model also uses communication intensity (currently a placeholder for future reference, e.g. number of function calls).

b) *PyAllocator* – a Python script which employs our multi-objective heuristic allocation method. Python was chosen for its simplicity in handling complex algorithms with rich and fast libraries (e.g. NumPy, SciPy). From a model created in Eclipse, SCALL extracts previously described matrices which are necessary for making an allocation decision. Once all the matrices are collected, they are converted to JSON format and sent to the PyAllocator which performs the allocation. Using AHP, PyAllocator calculates weights for each resource (trade-off vector \mathcal{F}) and verifies the consistency of the users input. Finally, the solution is obtained by a genetic algorithm and sent to Eclipse based model editor to be displayed for the user. Detailed description of the decision model can be found in [11].

4.1 Features

To provide an insight about the software allocation on a heterogeneous platform and also to give a performance estimation, SCALL currently has the following features:

a) *Model creation* — the user can simultaneously model software and hardware architecture of the system. For software components, the user can add resources and model their communication channels. The same can be done for computing units, within the same view.

b) *Model visualization* — enables easier model creation and better overview of the system. Instead of a classical approach of adding elements to a tree-like structure this approach enables drag-n-drop of model elements from a palette and a side-by-side view of both software and hardware architecture.

c) *Pairwise resource comparison* — Having a complete model ready for analysis, the user needs to perform a manual pairwise resource comparison using AHP notation [7] (i.e. determine which resources are more important).

d) *Software component allocation* — SCALL uses the described process to perform a multi-objective allocation of software components on a heterogeneous platform. The procedure results in a (sub-)optimal solution. At any time, the user can re-run the allocation process to get alternative allocations.

We generated architectures with different number of components and computing units with solution space varying between $[5^3, 15^{10}]$. SCALL performed very well, i.e. the average time necessary for the search spaces was (respectively) between $[0.59, 1.07]$ seconds while the exact solution found by full space search was up to 193 seconds². Given a genetic algorithm in the background, the exact solution was found for all tests and maximal error was 13% while, in average, SCALL was off by 3%.

4.2 Tool usage scenario

The following steps describe a typical usage scenario: a) determine software components, b) create a software architecture layout, c) determine computing units, d) create a hardware architecture layout, e) determine the necessary re-

sources for all software components, f) determine the available resources on each computing unit, g) perform a pairwise comparison of resources and h) find the satisfactory allocation.

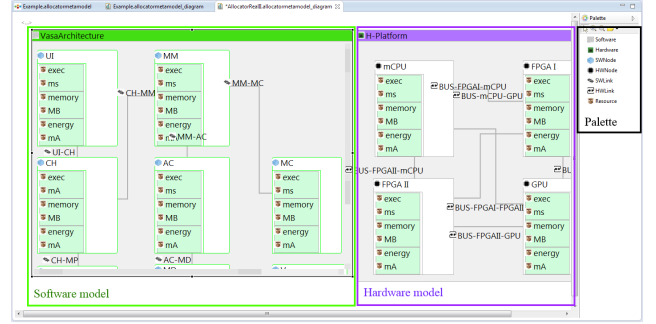


Figure 2: SCALL screen shot, showing software (left) and hardware (right) architecture side-by-side

In steps a) and c) the system architect needs to carefully consider functional and extra-functional properties to define components, interface contracts, computing units and communication channels. Links between components and computing units should be determined in steps b) and d) with consideration to bandwidth limitation. Steps e) and f) ensure the quantification of resources in a determined measurement unit. All the same resources need to be quantified in same measurement unit (e.g. sample model in Figure 2 considers memory in MB, computing time in μs and energy in W). For step g) the user needs to evaluate the significance of resources by a pairwise comparison. Higher mark for a resource gives more significance in allocation decision. In the final step, h) the user initiates the allocation process which provides a user with the results.

4.3 Deployment and installation

All the described features of SCALL along with the software dependencies are deployed as an Eclipse plugin. The installation can be found at the web page: <http://thor.foi.hr/scall>. Currently, SCALL is available only for Windows.

5. RELATED WORK

Due to high dimensionality of modern enterprise systems and adaptive software systems which emerge from ubiquitous computing, some authors [4] argue that there is a lack of formal methods to cope with the complexity which result in deployment alternatives. This is also very important in early design phase because in that phase mistakes are the cheapest and the notion of performance can guide a system design. Authors mostly focus on two parameters which influence the system performance; communication delay [3] and computational delay [8, 13].

A tool by Malek et al. addresses architecture optimization problem by assessing and improving the quality of deployment in a particular scenario or class of scenarios. Having scenario based decision they can multiple possibly conflicting QoS dimensions [5]. Their tradeoff model utilizes users' preferences for the desired levels of QoS, find the most suitable deployment architecture.

In [3] authors present an allocation method to improve re-

²2x Intel®Xeon®CPU E7-4830, 8GB Ram, Linux

source utilization and scheduling while keeping the decisions independent of platform implementation. The method determines the allocation using a genetic algorithm with an optimization function (heuristic algorithms are often used since allocation is a NP hard problem [9]).

A recent systematic review by Aleti et al. [2] shows that for optimizing software architecture researchers prefer to use custom modeling languages with allocation as a primary degree of freedom, followed by hardware replication and scheduling. However, current work largely focuses on allocating components and optimizing performance disregarding the emerging heterogeneous platforms. In this paper we are targeting high performance heterogeneous systems with various processing elements, e.g. CPU, GPU, FPGA.

6. POTENTIAL IMPACT AND FUTURE WORK

Allocating software components on computing units is an important challenge for the design decision makers in complex systems like automotive, avionics, manufacturing, etc. Our tool has a potential impact on: a) run-time and dynamic architectures b) (self-)adaptive systems and c) HSCD decision making. For a)–c) SCALL contributes to solving the challenge of design space exploration by supplying (sub-) optimal design alternatives which system architects need for better decision making focusing on early design phase performance evaluation. In an early design phase the notion of system performance can be used to guide the system design towards an efficient utilization of resources. SCALL is currently an early prototype and will be extended in future. There are several pending tasks to implement in the near future: a) enable generating the application skeleton from the model, with performance parameters included as variables which will be measured on real systems, b) improve on theory to remove more constraints (e.g. cross-computing unit components), c) refine the background algorithms through simulating performance evaluation with measured performance in existing systems.

Acknowledgment

SCALL was partially supported by Swedish Foundation for Strategic Research grant (SSF) via project *RALF3*.

7. REFERENCES

- [1] A. Aleti, S. Bjornander, L. Grunske, and I. Meedeniya. ArcheOpterix: An extendable tool for architecture optimization of AADL models. *odel-Based Methodologies for Pervasive and Embedded Software, 2009. MOMPES '09. ICSE Workshop on*, pages 61–71, 2009.
- [2] A. Aleti, B. Buhnova, L. Grunske, A. Koziolok, and I. Meedeniya. Software Architecture Optimization Methods: A Systematic Literature Review. *IEEE Transactions on Software Engineering*, 39(5):658–683, May 2013.
- [3] T. Çelik and B. Tekinerdogan. S-IDE: A tool framework for optimizing deployment architecture of high level architecture based simulation systems. *Journal of Systems and Software*, 86(10):2520–2541, 2013.
- [4] L. Da Xu. Enterprise systems: State-of-the-art and future trends. *Industrial Informatics, IEEE Transactions on*, 7(4):630–640, 2011.
- [5] S. Malek, N. Medvidovic, and M. Mikic-Rakic. An extensible framework for improving a distributed software system’s deployment architecture. *IEEE Transactions on Software Engineering*, 38(1):73–100, 2012.
- [6] B. Ristau, T. Limberg, and G. Fettweis. A mapping framework for guided design space exploration of heterogeneous MP-SoCs. *Proceedings of the conference on Design, automation and test in Europe - DATE '08*, page 780, 2008.
- [7] T. Saaty. *Fundamentals of Decision Making and Priority Theory with the Analytic Hierarchy Process*. Rws Publications, 1994.
- [8] L. Santinelli, M. Chitnis, C. Nastasi, F. Checconi, G. Lipari, and P. Pagano. A component-based architecture for adaptive bandwidth allocation in wireless sensor networks. In *Industrial Embedded Systems (SIES), 2010 International Symposium on*, pages 174–183. IEEE, 2010.
- [9] A. A. Ssaed, W. Kadir, M. Wan, M. Hashim, and S. Zaiton. Metaheuristic search approach based on in-house/out-sourced strategy to solve redundancy allocation problem in component-based software systems. *International Journal of Software Engineering & Its Applications*, 6(4), 2012.
- [10] J. Teich. Hardware/software codesign: The past, the present, and predicting the future. *Proceedings of the IEEE*, 100(Special Centennial Issue):1411–1430, 2012.
- [11] I. Švogar, I. Crnković, and N. Vrček. An extended model for multi-criteria software component allocation on a heterogeneous embedded platform. *Journal of Computing and Information Technology*, 21(4):211–222, 2013.
- [12] V. Vyatkin. Software engineering in industrial automation: State-of-the-art review. *Industrial Informatics, IEEE Transactions on*, 9(3):1234–1249, Aug 2013.
- [13] S. Wang and K. G. Shin. Early-stage performance modeling and its application for integrated embedded control software design. In *ACM SIGSOFT Software Engineering Notes*, volume 29, pages 110–114. ACM, 2004.

APPENDIX

Demonstration plan:

- *Introduction and motivation* — explain the issues related to software architecture planning and early design phases
- *Background* — briefly present the problem background and explain why component allocation on heterogeneous systems is an important problem.
- *Live tool demo* — we will apply the tool on a concrete problem related to allocation of software components on a heterogeneous system. In particular, it's an underwater autonomous vehicle (AUV) with 3 main computing units, *multicore CPU*, *GPU* and *FPGA* with 11 software components. First few components will be added manually, and the final example will be loaded from file.
- *Conclusion and future work* — conclude the presentation with current tool drawbacks, future implementation and potential impact.
- *Q&A*

Tool availability and maturity

SCALL is currently a prototype with some pending updates to be implemented in the near future. These include code a) skeleton generation, b) model update with additional optimization algorithms (simulated annealing), c) an alternative to AHP and d) re-validation of the underlying mathematical model in a real system (a robot with a heterogeneous platform, with different software allocations competing). A tool manual will also be available in the future, as suggested by users. Unfortunately, the tool is currently available only for Windows:

URL: <http://thor.foi.hr/scall>

CATEGORY: research

VIDEO: <https://youtu.be/dbsQZ36susw>